

An Evaluation of Protocols for UAV Science Applications

William D. Ivancic
NASA Glenn Research Center

David E. Stewart
Verizon Federal Systems

Donald V. Sullivan
Patrick E. Finch
NASA Ames Research Center

Abstract— This paper identifies data transport needs for current and future science payloads deployed on the NASA Global Hawk Unmanned Aeronautical Vehicle (UAV). The NASA Global Hawk communication system and operational constraints are presented. The Genesis and Rapid Intensification Processes (GRIP) mission is used to provide the baseline communication requirements as a variety of payloads were utilized in this mission. User needs and desires are addressed. Protocols are matched to the payload needs and an evaluation of various techniques and tradeoffs are presented. Such techniques include utilization rate-base selective negative acknowledgment protocols and possible use of protocol enhancing proxies. Tradeoffs of communication architectures that address ease-of-use and security considerations are also presented.

TABLE OF CONTENTS

I. BACKGROUND	1
II. COMMUNICATION ARCHITECTURE	1
III. REQUIREMENTS	3
IV. TRANSPORT PROTOCOLS	3
V. PROTOCOL ENHANCING PROXIES	5
VI. CONCLUSIONS	5
REFERENCES	6

I. BACKGROUND

NASA has acquired two Global Hawk Unmanned Aerial Vehicles (UAVs) to enhance upper atmospheric science in support of the Earth Science Project Office (ESPO). There is a potential to acquire an additional 3 to 5 Global Hawks resulting in a fleet of 5 to 7 aircraft. The Global Hawks currently reside at the Dryden Research Center (DRC). DRC is located on Edwards Air Force Base. Edwards has one of the largest military “special use” areas in the United States covering almost 16,000 square miles. “Special Use” airspace is required to get the Global Hawks from ground to evaluations above commercial air traffic.

The Global Hawk has a maximum endurance of 42 hours with an on-station endurance of 24 hours at 3,000 Nm from point of departure. It can loiter at 343 knots and has a maximum altitude of 65,000 feet, which is above the weather and above the commercial air space. Such capabilities allow the Global Hawk to take unique measurements at a fixed area over a full solar day. As

such, the Global Hawk measurements are complimentary to satellite data.

II. COMMUNICATION ARCHITECTURE

Security

For the NASA operated Global Hawks, the aircraft operation’s command and control communications is completely separate from the experimental payloads’ command and control. This enables different security methodologies to be deployed for each system. Aircraft control is critical for safety of flight and safety of life. Furthermore, hijacking a UAV basically puts a missile in the hands of the hijacker. By separating the command of the payload from the command of the aircraft, the security required for payload operations becomes much less stringent and enables greater flexibility of payload deployment and direct real-time access to payload instrumentation by the various principle investigators.

To date, securing the payload has been accomplished using standard computer access techniques. User accounts are established for each Principal Investigator that needs payload access. Login is via Secure Shell (SSH) or telnet. There is currently no requirement in the payload communication chain to secure the RF link or for network layer security, Internet Protocol Security (IPsec). If network layer security is required in future missions, IPsec could be easily enabled.

Satellite Communications

Communication with the experimental payload is by a Ku-Band satellite link. Initial deployments used a 2 Mbps bidirectional link. Future flights are expected to use up to 8 Mbps links. The system is capable of approximately 50 Mbps but the cost to operate at such rates is prohibitive. Furthermore, there currently is not a requirement to move the volume of data that would require a 50 Mbps link although there may be in the future.

During the Global Hawk Pacific (GloPac) mission, Ku-Band connectivity was demonstrated to approximately 75 degrees north latitude (approximated 3 degrees elevation angle). Thus, any missions with flight profiles below 75 degrees (north or south) have continuous connectivity.

In addition to the Ku-band satellite link, four Iridium L-Band modems have been multiplexed together to provide some low-rate (kbps) communication to the experimental payloads for simple commanding and status (telemetry).

Deployment Scenarios

The vast majority of payload communication and control available to the principle investigators is via the Ku-band satellite links. The Ku-band satellite links are the only links with sufficient bandwidth to warrant any protocol related data transport optimization. Therefore, the remainder of this paper will only be concerned with high-bandwidth geostationary satellite links when considering communication between the Global Hawk payloads and the Principle Investigators.

Figures 1a, 1b and 1c illustrate three possible deployment scenarios for Ku-Band satellite communication. Figure 1a is the current communication deployment. A Ku-band terminal is located at Dryden Research Center (DRC) with a fiber optic run between the ground station and the control center. Communication from PI to Payload is effectively a direct link resulting in a single communication control loop. This was the scenario used in the GloPac and Genesis and Rapid Intensification




Figure 1a - Current Communication Deployment




Figure 1b - Venture Mission (Atlantic Campaign)




Figure 1c - Future Deployment Possibilities

Processes (GRIP) missions. For GloPac, RF communication over the Ku-band link was lost over the North Pole region at which point only Iridium links could be used to check payload status. No large volumes of data could be downloaded through the Iridium links due to extremely small bandwidths. For GRIP, RF communication over the Ku-band link was continuous as the mission was over the tropics with the satellite in constant contact with the Global Hawk.

The only difference between scenario 1a and 1b is that a transportable ground station will be deployed in close enough proximity to the mission rendezvous point to enable constant contact between the Global Hawk, the Ku-band satellite and the transportable ground station. Since the Principle Investigators will be collocated at the transportable ground station site, the optimal communication protocol choice is identical for scenarios 1a and 1b.

Future deployments may utilize the scenario presented in figure 1c. There is even a possibility that such a scenario may be used for the Earth Venture Mission - Hurricane and Severe Storm Sentinel (HS3) campaign [1]. Here, the RF modem and networking equipment would be located at the remote ground station with the Principle Investigators located at, and performing operations from DRC. In such a scenario, two major architectural deployments are possible. The first architecture would have the initial mission data-storage taking place directly at the ground station. In this case, there is only one control loop and everything is identical to scenarios 1a and 1b. The PIs could login to the remote data storage computer to execute payload commanding. The second architecture would have the initial mission data storage at DRC. That would result in either a single long control loop over both the space /ground link and the ground/ground link or a situation where one would break the control loops between the space/ground link and the ground/ground link. The latter would lend itself to deployment of store and forward techniques such as delay tolerant networking (DTN), a network overlay, or other much simpler store and forward techniques.

Communication Network Characteristics

In all three scenarios, there is a single ground station and a single communication path from PI to payload. Therefore, from a networking perspective, the Global Hawk is stationary and solutions to address network mobility are not required.

The Ku-band geostationary satellites have approximately 500 milliseconds of route trip time (RTT) delay. Allowing for additional processing and queuing may result in up to 600 milliseconds of RTT delay.

The current modems and antennas used on the Global Hawks provide near-error-free links – particularly at the data rates used (less than 10 Mbps).

III. REQUIREMENTS

Protocol Requirements

The primary requirement of the protocols is to provide a good user experience while remaining as indistinguishable as possible from existing Internet protocols. Part of a good user experience is getting the required science data down in a timely manner. This has to be done while operating over near-error-free links with 600 milliseconds RTT delays.

User Requirements

The users of this system are the scientists (i.e. the Principal Investigators and their collaborators). This group is interested in ease of use and maximum delivery of science data. Their preference is to use as many existing Internet protocols as possible. Doing so allows the scientists to test their instruments and data collecting in the lab, on the ground, and in flight using the same protocols, commands, and scripts. The PIs desire to use the exact same Internet tools used in the lab while operating on the DC-8 research aircraft or while controlling instrumentation onboard the Global Hawk. Note: on the DC-8 the scientist is collocated with the payload effectively making the RTT delay a few milliseconds, whereas, when controlling instrumentation on the Global Hawk, the RTT is approximately 600 milliseconds.

The protocols currently used are all based on the Transmission Control Protocol (TCP). They include: Telnet, Secure Shell (SSH), and file transfer protocols (i.e. File Transfer Protocol (FTP), Secure Copy Protocol (SCP), Secure File Transfer Protocol (SFTP), RSYNC, WGET, etc...). Often the file transfer protocols are run in an SSH tunnel.

Research Data Requirements

The Lightning Instrumentation Package (LIP) measures lightning, electric fields, electric field changes, and air conductivity. The amount of data produce is relatively small. The data throughput requirement is kbps. Such low telemetry needs can be met by Iridium.

The High Altitude MMIC Sounding Radiometer (HAMSR) was developed by JPL. It operates using 25 spectral channels in 3 bands and provides measurements that can be used to infer the 3-D distribution of temperature, water vapor, and cloud liquid water in the atmosphere. The data requirements are approximately 200 Mbytes over duration of mission (24 hours) with instantaneous throughputs of 10s to 100s of kbps. The current system uses RSYNC over TCP to synchronize the

ground database with payload database. RSYNC is run periodically during periods when the Global Hawk is in contact with the ground station. This is a very simple and effective store and forward technique that works well for the operational scenarios presented. Telnet or SSH are used to login and periodically check the instrument status.

The High-Altitude Imaging Wind and Rain Airborne Profiler (HIWRAP) is a dual-frequency radar (Ka- and Ku-band), dual-beam (300 and 400 incidence angle), conical scan, solid-state transmitter-based system, designed for operation on the high-altitude (20 km). By combining measurements at Ku- and Ka-band, HIWRAP is able to image winds by measuring volume backscattering from clouds and precipitation. Operators use telnet or SSH to check payload status.

For the GRIP flights, the digital receiver was functioning and obtaining raw data by recording in-phase and quadrature-phase (I and Q) of the return signal. This resulted in a data rate of about 1 Gigabyte per minute (approximately 130 Mbps), which exceeds the available links even at continuous operation. A 1 Terabyte disk was filled during flight. This data was then offloaded after the Global Hawk landed. During GRIP, there was no processing that averaged, compressed, and produced science parameters. By deploying such onboard processing on future flights, the data-rate should be reduced by a factor of about 15, or 66 MB per minute (8.8 Mbps link requirement).

Another approach being considered is to produce real-time data products by using FPGA-based processing. Quicklook products such as images would be produced rather than sending back raw data. These techniques would greatly reduce the data downlink requirements to well within the current bandwidth of the Ku-band communication system.

IV. TRANSPORT PROTOCOLS

TCP vs. UDP Rate-Based

There are two basic types of transport protocols: the Transmission Control Protocol (TCP) and transmission protocols built at the application layer using the User Datagram Protocol (UDP) as the network transport mechanism.

TCP Operations

Most Internet applications use TCP as the underlying transport protocol. TCP is a reliable transport protocol built into the kernel of computer operating systems. It guarantees reliable, ordered delivery of a stream of bytes from a program on one computer to another program on another. TCP is designed to operate over shared networks. It includes flow control and congestion control algorithms that work well in the terrestrial Internet. The

congestion control algorithm design is such that operation over noisy and long delay links is problematic. TCP will guarantee data delivery, but the link will not be used efficiently as TCP assumes any loss to be caused by congestion.

TCP initially aggressively probes the network via an exponential increase in data-rate. TCP will double its data-rate each round trip time (RTT) by doubling the amount of packets sent (e.g. 4, 8, 16, 33) until it senses network congestion. As soon as TCP senses a packet loss (network congestion), it cuts its data-rate in half and then conservatively probes the network by increasing its transmission rate one packet per round trip time (RTT). This conservative probing is called “linear increase”. This combination of algorithms is excellent for a shared network and relatively low delays. However, for a private link where no congestion is present and where errors may be present, these algorithms have problems reaching and maintaining maximum link capacity. Furthermore, unless one manually tunes the flow control sliding-window parameters, TCP will eventually self-congest and cut its data-rate in half resulting in the saw-tooth transmission pattern illustrated in figure 2. Note: tuning protocols is non-trivial.

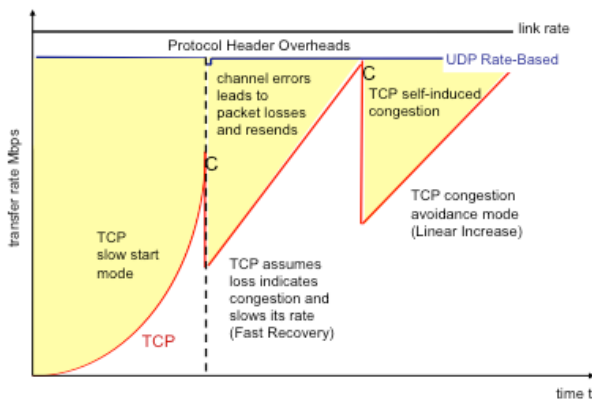


Figure 2 TCP vs. UDP Rate-Base Performance

UDP Rate-base Operations

UDP rate-base protocols operate at line-rate or at some set rate-limit. How one determines the rate-limit is deployment and implementation specific. UDP rate-based protocols generally assume no congestion and thus deploy no congestion control algorithms. Thus there is no need to probe the system to determine available bandwidth or to reduce data-rates when losses occur as all losses are assumed to be due to errors rather than congestion. UDP-based transport protocols utilize a negative acknowledgement algorithm (NACK) to notify the sender when errors occur. Here, the receiver explicitly notifies the sender of which packets, messages, or segments were received incorrectly and thus may need to be

retransmitted. Figure 2 illustrates that UDP-based transport protocols provide throughput near the line-rate of the transmission media (minus the protocol overhead). Figure 3 shows the throughput of a rate-base NACK protocol versus TCP for large file transfers. This graph is for a packet size of 1024 bytes and assumes a binomial error distribution. The TCP portion of the graph is derived from a well-documented TCP performance equation [2]. Note: data throughput for a rate-based protocol is independent of propagation delay (RTT) whereas TCP performance is heavily affected by RTT.




Figure 3 Theoretical Throughput of TCP vs. Rate-Based Protocols for 1024 byte packets

A non-comprehensive list of UDP-based transport protocols includes: Saratoga, Negative Acknowledgement (NACK) - Oriented Reliable Multicast (NORM), Licklider Transmission Protocol (LTP), and the Consultative Committee for Space Data Systems (CCSDS) File Delivery Protocol (CFDP). All of these use a selective, negative acknowledgment mechanism for transport reliability.

“Saratoga is a simple, lightweight, content dissemination protocol that builds on UDP, and optionally uses UDP-Lite. Saratoga is intended for use when moving files or streaming data between peers, which may have permanent, sporadic or intermittent connectivity, and is capable of transferring very large amounts of data reliably under adverse conditions.” Surrey Satellite Technology Limited (SSTL) originally implemented Saratoga to efficiently transfer remote-sensing image from a low-Earth-orbiting satellite to ground over highly asymmetric links. It has been used in SSTL’s Disaster Monitoring Constellation satellites since 2004 [3].

“Saratoga also has a beacon that can be activated. The beacon is used to announce the presence of the node to potential peers (e.g. satellites, ground stations) as well as to provide automatic service discovery, and to confirm the activity or presence of the peer.”

“The primary design goals of NORM are to provide efficient, scalable, and robust bulk data (e.g., computer files, transmission of persistent data) transfer across

possibly heterogeneous IP networks and topologies. The NORM protocol design provides support for distributed multicast session participation with minimal coordination among senders and receivers. NORM allows senders and receivers to dynamically join and leave multicast sessions at will.” NORM leverages the use of forward error correction (FEC) repair and other IETF Reliable Multicast Transport (RMT) building blocks. NORM also has a congestion control scheme to fairly share available network bandwidth with other transport protocols such as TCP [4]. In addition, NORM can operate in unicast mode as had been demonstrated on the Naval Research Lab’s MidStar-1 Satellite for unidirectional link file transfer [5].

“LTP and CFDP are designed to provide retransmission-based reliability over links characterized by extremely long message round-trip times (RTTs) and/or frequent interruptions in connectivity such as found in deep-space communications.” LTP and CFDP are based in part on a notion of massive state retention, which is necessary over extremely long delays in order to keep track of acknowledgments for data sent. LPT and CFDP have retransmission timers that must be suspended during periods of disconnection. This can be done either via per-configuration on known link outages due to orbital dynamics and/or by passing information on the link status to the LTP protocol engine (link state queues) [6,7,8].

NORM, LTP and CFDP are rather complex protocols with capabilities that are not required in the Global Hawk payload system. NORM is intended for multicast enabled systems and large groups of receivers. In addition, NORM has FEC capability that is not necessary for our near-error-free communication links. LTP and CFDP are designed for extremely long delays and disconnection – neither which are found in our communication system. LTP and CFDP also require configuration of many parameters in order to tune the protocol to the particular environment it is being used in. Saratoga has its roots in CFDP and is actually a scaled-down, simplified version of CFDP. We anticipate Saratoga to be the protocol of choice for large file transfers due to its simplicity, ease of use, and built-in discovery features.

V. PROTOCOL ENHANCING PROXIES

Protocol Enhancing Proxies (PEPs) are used to improve TCP performance over long delays. The basic idea is to break the end-to-end control loop into multiple control loops such that one can utilize a protocol that performs well over long-delay, error prone links without modifications to the end users system (protocols). In this manner, the end users can use standard TCP-based Internet protocols without experiencing performance degradations. However, PEPs have known problems. They require a reasonable amount of additional processing, often require special configuration and tuning – particularly with regard to the available bandwidth and

RTT delay – and require special care in deployment when used in conjunction with IPsec [9].

In the case of the Global Hawk, the delay is approximately 600 milliseconds with near-error-free performance. The PEP would utilize a rate-based protocol between the ground station and Global Hawk (control loop 2) while standard TCP would be run at on the Principal Investigators’ systems and on the Payload Control system (control loops 1 and 3) [Figure 4].




Figure 4 Protocol Enhancing Proxy Deployments

By deploying PEPs, one should expect improvements in large file transfers assuming standard File Transfer Protocol (FTP) or Secure Copy Protocol (SCP) are used. Such data delivery improvements should be comparable to UDP rate-based file transfer protocols such as Saratoga or NORM. For very small files or command and control messages, one should experience very little improvement, with a PEP compared to having no PEP in the system. Note: PEPs will not help interactive communications, as PEPs cannot remove the propagation delay.

VI. CONCLUSIONS

During the GloPac and GRIP missions, the Global Hawk payloads were directly accessible to the Principle Investigators using standard Internet protocols with no PEPs deployed. The user experience was positive even without PEPs. This is primarily because the users only accessed their payloads to obtain status or to configure their systems. Such communication requires very small file transfers or simple message exchanges. Larger file transfers for GRIP and GloPac were performed in the background using RSYNCH for remote synchronization. As such, any TCP inefficiencies were not apparent to the user.

In the future deployments, real-time delivery of larger data will be required and efficient use of the communication links will be necessary. At that time, either PEPs or an efficient, rate-based protocol such as Saratoga or both will be installed depending on the performance needs are architectural deployment¹. Use of only a rate-based protocol is preferred over deployment of PEPs in order to keep the communication system as simple as possible.

¹ At the time of publication, performance tests were being performed to evaluate the CCSDS Space Communication Protocol Suit (SCPS) PEP relative to UDP-based file transfer protocols. The results of that testing are expected to be formally published in a follow-on document.

REFERENCES

- ¹ “Hurricane and Severe Storm Sentinel”, Earth System Science Pathfinder Mission
<http://science.nasa.gov/missions/hs3/>
- ² Mathis, M. et al, "The Macroscopic Behavior of the Congestion Avoidance Algorithm," Computer Communications Review, volume 27, number 3, July 1997.
- ³ L. Wood, W. Eddy, C. Smith, W. Ivancic, C. Jackson: “Saratoga: A Scalable File Transfer Protocol,” draft-wood-tsvwg-saratoga-08, December 23, 2010 (Work in Progress)
- ⁴ C. Bormann, M. Handley, J. Macker: “Naval Research Laboratory NACK-Oriented Reliable Multicast (NORM) Transport Protocol,” IETF RFC 5740, November 2009
- ⁵ K. Hogue, “NORM with requeue now operational on MidSTAR-1,” email to IETF norm-dev mail list, December 17, 2008
<http://pf.itd.nrl.navy.mil/pipermail/norm-dev/2008-December/000554.html>
- ⁶ S. Burleigh et al., Licklider Transmission Protocol - Motivation, IETF RFC 5325, informational September 2008
- ⁷ M. Ramadas et al., Licklider Transmission Protocol - Specification, IETF RFC 5326, experimental, September 2008
- ⁸ CCSDS File Delivery Protocol, Issue 2, CCSDS 727.0-B-2, October 2002
- ⁹ J. Border, M. Kojo, J. Griner, G. Montenegro, Z. Shelby: “Performance Enhancing Proxies Intended to Mitigate Link-Related Degradations,” IETF RFC 3135, June 2001